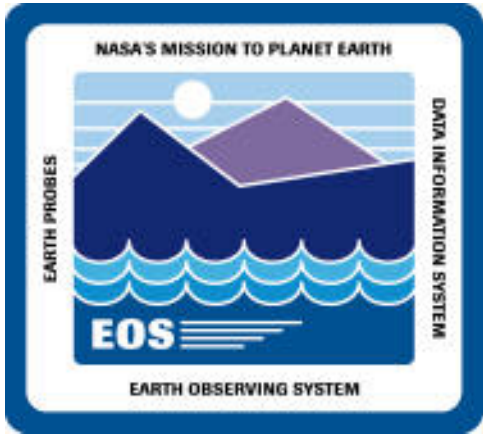

DPS HW Design Review

Randy Miller

randallm@eos.hitc.com

19 April 1996

Overview

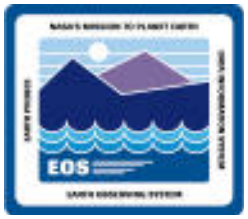

- **Data Processing Subsystem (DPS) Hardware Configuration Items**
- **Requirements**
- **Sizing Analysis**
- **Specification**
- **Design Analysis**
- **Design Validation**

DPS Hardware Configuration Items

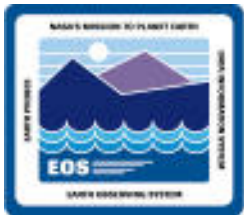

- SPRHW (Science Processing Hardware CI)
  - Compute resources for PGEs and AI&T
  - Includes Queuing Server (to be discussed in Planning Subsystem presentation)
- AITHW (Algorithm Integration and Test Hardware CI)
  - Server with tools for AI&T
  - Seats for AI&T personnel
- AQAHW (Algorithm Quality Assurance Hardware CI)
  - Visualization workstation for DAAC-based non-science QA

DPS Hardware Diagram

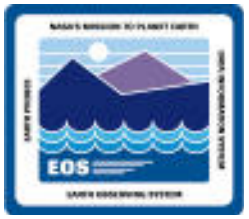

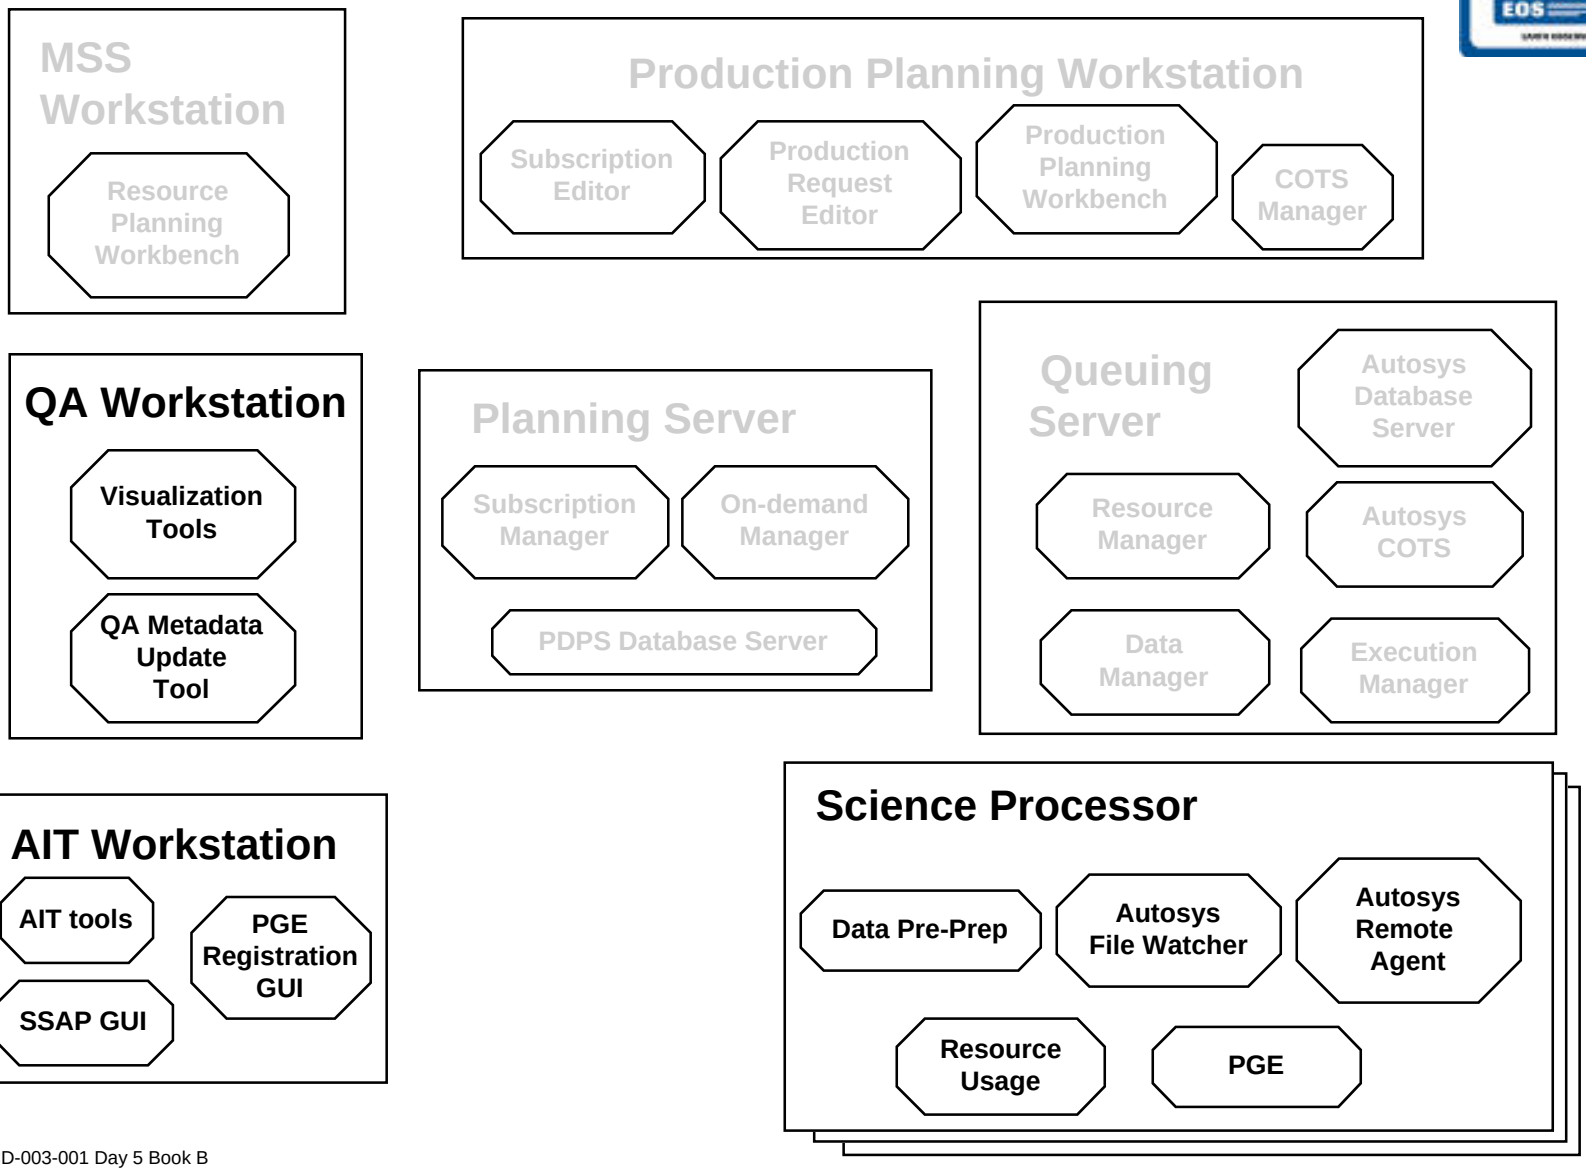

SPRHW F&PRS Requirements

- Timeliness
  - EOSD1050, EOSD1060, EOSD1070: Level 1, Level 2, and Level 3 products must be generated and available within 24 hours of input data availability
    - Production of certain MODIS Level 3 products within 24 hours of data availability causes a high ratio of peak to average processing requirements (10:1) at EDC
    - A CCR is in progress to modify this requirement so that these products can be produced over a longer time span (approximately five days)
    - The CDR design assumes that the CCR will be accepted, and sizes EDC accordingly
  - EOSD1080: Level 4 products must be generated and available within 7 days of input data availability

SPRHW F&PRS Requirements

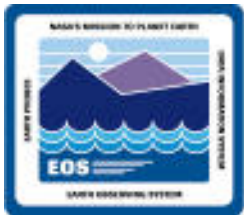

- **Allocation of Resources**
  - **EOSD1040: Reprocessing capacity shall be twice the first-time processing capacity**
  - **PGS-1300: A capacity for "algorithm and test demands, production of prototype products, ad hoc processing for 'dynamic browse' or new search and access techniques developed by science users, and additional loads due to spacecraft overlap" shall be provided equal to the first-time processing capacity**
- **Expandability**
  - **PGS1270: Design must accommodate expansion by up to a factor of 3 without design changes, and up to a factor of 10 without major design changes**
- **Derating of Processor Performance**
  - **PGS-1301: Processing rates used for sizing shall not be greater than 25% of peak-related CPU capacity**

SPRHW Phasing Requirements

Procurement of SPRHW shall be phased based upon the launch dates of the instruments:

- **<u>0.3X for L-2 < t < L-1</u>. Pre-launch AI&T requires 0.3X during the period from one to two years before launch, where X is defined as the resource requirement for first-time processing of instrument data.**
- **<u>1.2X for L-1 < t < L+1</u>. Pre-launch AI&T and system I&T requires 1.2X during the year before launch. First-time processing requirements (X) begin from the launch date and last for the remainder of the life of the instrument.**
- **<u>2.2X for L+1 < t < L+2</u>. Post-launch AI&T, standard processing, and reprocessing of data require 2.2X starting at launch plus one year.**
- **<u>4.2X for t > L+2</u>. Post-launch AI&T, standard processing, and reprocessing of data require 4.2X starting at launch plus two years.**

SPRHW Phasing Requirements

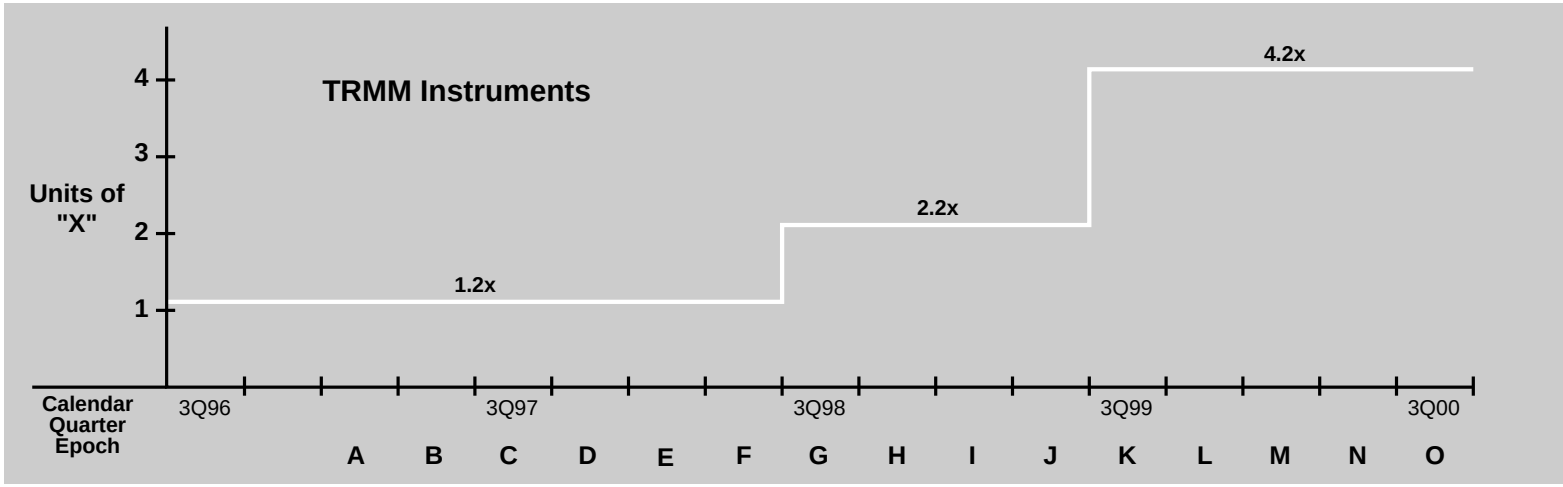

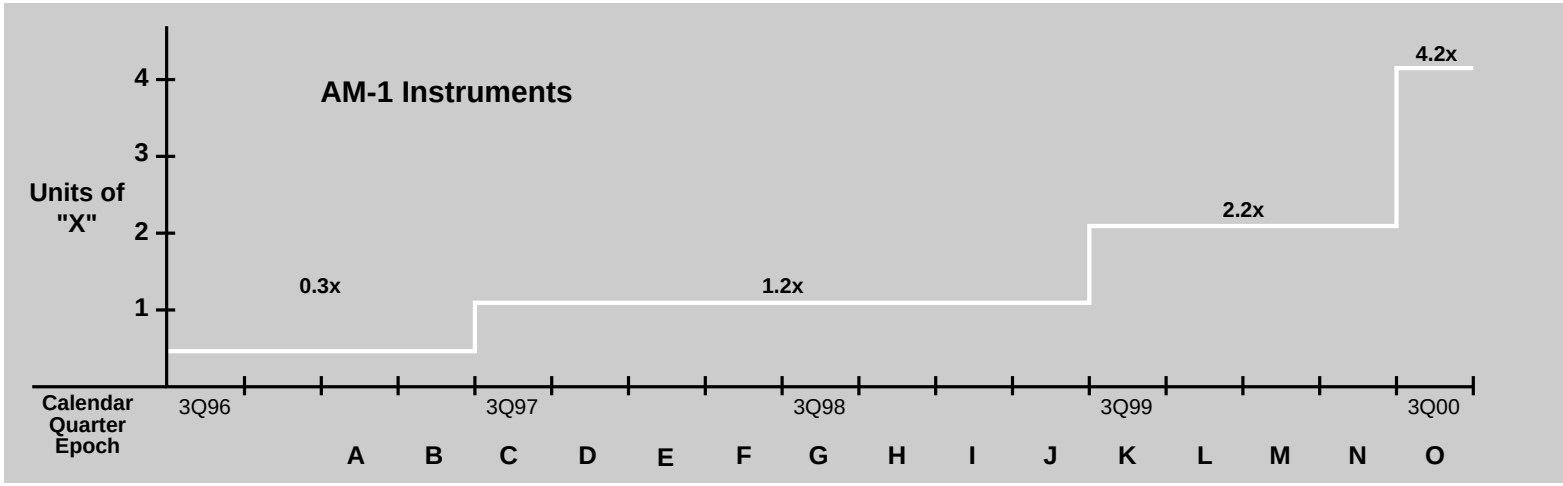

Other SPRHW Requirements

- Selected hardware must support standard ECS software
- Selected hardware must support highspeed interconnect protocol
  - HiPPI, OC-12 ATM, or FC-AL
- RMA requirements
  - Availability of 96% or better
  - Mean down time not greater than 4 hours

Sizing Approach

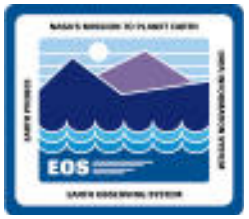

- Static modeling of AHWGP inputs
  - Provides average loads for CPU, networks, disk I/O
- Dynamic modeling
  - Provides more accurate loads for CPU, networks, disk I/O, and disk size
- Memory and I/O survey
  - To provide memory requirements and refinements in I/O modeling
- Other considerations and adjustments

Static Modeling Analysis

- **Process descriptions and volume timelines from the Ad Hoc Working Group for Production (AHWGP) are entered into a spreadsheet**

- **Processing requirements (CPU, network I/O, disk I/O, archive I/O) are summed by instrument and DAAC**

- **This approach supports several analyses:**
  - **Average load over long periods of time**
  - **Peak load on the worst case day**
  - **Loads under other assumptions (for example, spreading Level 3 production over multiple days)**

Static Modeling Results (LaRC)

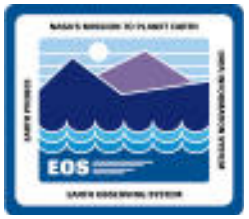

Process ID	Instrument	Volume at Initiation (MB)	Staging I/O (MB)	Volume at Completion (MB)	Destaging I/O (MB)	Total I/O (MB)	Millions of Floating Point Ops per Execution	# Read per Execution	# Written per Execution	No. of Exec. /day	Processing (MFLOPS)	I/O Local to Processing (MB/sec)	Worst Case Processing I/O (MB/sec)	Worst Case Network I/O (MB/sec)	Best Case Deep Arch I/O (MB/sec)
10aA	CERES(AM)	25,791	25,785	26,920	1,129	26,920	245,700	896.0	1.0	0.03	0.1	0.0	0.0	0.0	0.0
1aA	CERES(AM)	99	92	1,536	1,431	1,536	20,790	9.0	28.0	1.00	0.2	0.0	0.0	0.0	0.0
1bA	CERES(AM)	99	92	1,536	1,431	1,536	20,790	9.0	28.0	1.00	0.2	0.0	0.0	0.0	0.0
2aA	CERES(AM)	627	627	903	174	804	3,780	5.0	3.0	1.00	0.0	0.0	0.0	0.0	0.0
2bA	CERES(AM)	102	0	203	102	105	1	1.0	1.0	0.03	0.0	0.0	0.0	0.0	0.0
3aA	CERES(AM)	102	102	676	574	676	47,250	3.0	4.0	0.03	0.0	0.0	0.0	0.0	0.0
4bAF	CERES(AM)	371	351	11,467	11,096	11,467	3,440,000	14.0	5.0	24.00	955.6	3.2	6.4	3.2	3.1
5cAF	CERES(AM)	463	433	902	439	902	2,672,460	7.0	1.0	24.00	742.4	0.3	0.5	0.2	0.1
5cAV	CERES(AM)	463	433	902	439	902	2,672,460	7.0	1.0	4.00	123.7	0.0	0.1	0.0	0.0
6aA	CERES(AM)	463	461	493	22	493	4,914	4.0	2.0	24.00	1.4	0.1	0.3	0.1	0.0
6cA	CERES(AM)	6,212	0	12,423	6,211	12,423	4	744.0	144.0	0.03	0.0	0.0	0.0	0.0	0.0
7aA	CERES(AM)	29,039	28,949	45,461	16,368	45,461	4,082,400	901.0	249.0	0.03	1.5	0.0	0.0	0.0	0.0
8aA	CERES(AM)	16,371	16,368	17,104	733	17,104	226,800	251.0	2.0	0.03	0.1	0.0	0.0	0.0	0.0
9aAF	CERES(AM)	342	340	361	16	361	4,914	4.0	2.0	24.00	1.4	0.1	0.2	0.1	0.0
9bAF	CERES(AM)	3,125	0	6,250	3,125	6,250	4	744.0	144.0	0.03	0.0	0.0	0.0	0.0	0.0
11	CERES(TRMM)	6,200	6,200	6,540	340	6,540	12,600	1240.0	1.0	0.03	0.0	0.0	0.0	0.0	0.0
10aT	CERES(TRMM)	25,791	25,785	26,920	1,129	26,920	245,700	896.0	1.0	0.03	0.1	0.0	0.0	0.0	0.0
10bTA	CERES(TRMM)	28,916	28,910	30,045	1,129	30,045	491,400	1040.0	1.0	0.03	0.2	0.0	0.0	0.0	0.0
12aF	CERES(TRMM)	60	60	780	720	780	37,800	10.0	24.0	1.00	0.4	0.0	0.0	0.0	0.0
1aT	CERES(TRMM)	256	250	1,694	1,431	1,694	20,790	10.0	28.0	1.00	0.2	0.0	0.0	0.0	0.0
2aT	CERES(TRMM)	627	627	903	174	804	3,780	5.0	3.0	1.00	0.0	0.0	0.0	0.0	0.0
2bT	CERES(TRMM)	102	0	203	102	105	1	1.0	1.0	0.03	0.0	0.0	0.0	0.0	0.0
3aT	CERES(TRMM)	102	102	676	574	676	47,250	3.0	4.0	0.03	0.0	0.0	0.0	0.0	0.0
3bTA	CERES(TRMM)	203	203	778	574	778	94,500	4.0	4.0	0.03	0.0	0.0	0.0	0.0	0.0
4aF	CERES(TRMM)	114	114	3,096	2,982	3,069	72,000	4.0	4.0	24.00	20.0	0.9	1.7	0.9	0.8
5aF	CERES(TRMM)	463	433	902	439	902	2,672,460	7.0	1.0	24.00	742.4	0.3	0.5	0.2	0.1
5aV	CERES(TRMM)	463	433	902	439	902	2,672,460	7.0	1.0	4.00	123.7	0.0	0.1	0.0	0.0
6aT	CERES(TRMM)	463	461	493	22	493	4,914	4.0	2.0	24.00	1.4	0.1	0.3	0.1	0.0
6cT	CERES(TRMM)	6,212	0	12,423	6,211	12,423	4	744.0	144.0	0.03	0.0	0.0	0.0	0.0	0.0
7aT	CERES(TRMM)	29,039	28,949	45,461	16,368	45,461	4,082,400	901.0	249.0	0.03	1.5	0.0	0.0	0.0	0.0
7c	CERES(TRMM)	35,249	35,160	51,671	16,368	51,671	8,164,800	1045.0	249.0	0.03	3.0	0.0	0.0	0.0	0.0
8aT	CERES(TRMM)	16,371	16,368	17,104	733	17,104	226,800	251.0	2.0	0.03	0.1	0.0	0.0	0.0	0.0
8c	CERES(TRMM)	16,371	16,368	17,104	733	17,104	453,600	251.0	2.0	0.03	0.2	0.0	0.0	0.0	0.0
9aTF	CERES(TRMM)	342	340	361	16	361	4,914	4.0	2.0	24.00	1.4	0.1	0.2	0.1	0.0
9bTF	CERES(TRMM)	3,125	0	6,250	3,125	6,250	4	744.0	144.0	0.03	0.0	0.0	0.0	0.0	0.0
MISP1A-AA-1-0LM	MISR	212	212	763	550	763	29,398	5.0	5.0	0.31	0.1	0.0	0.0	0.0	0.0
MISP1A-AA-2-0LM	MISR	5,379	456	6,584	1,205	6,584	675,693	9.0	9.0	9.25	72.3	0.7	0.9	0.2	0.1
MISP1A-AA-2-1LM	MISR	5,384	461	6,604	1,220	6,604	676,637	9.0	10.0	4.00	31.3	0.3	0.4	0.1	0.1
MISP1A-AA-2-2LM	MISR	5,384	461	6,613	1,228	6,613	677,582	9.0	11.0	1.00	7.8	0.1	0.1	0.0	0.0
MISP1A-AF-1-0LM	MISR	212	212	763	550	763	29,398	5.0	5.0	0.31	0.1	0.0	0.0	0.0	0.0
MISP1A-AF-2-0LM	MISR	5,379	456	6,584	1,205	6,584	675,693	9.0	9.0	9.25	72.3	0.7	0.9	0.2	0.1
MISP1A-AF-2-1LM	MISR	5,384	461	6,604	1,220	6,604	676,637	9.0	10.0	4.00	31.3	0.3	0.4	0.1	0.1
MISP1A-AF-2-2LM	MISR	5,384	461	6,613	1,228	6,613	677,582	9.0	11.0	1.00	7.8	0.1	0.1	0.0	0.0

Static Modeling Results (Summary)

		No. of Exec. /day	Processing (MFLOPS)	I/O Local to Processing (MB/sec)	Worst Case Processing I/O (MB/sec)	Worst Case Network I/O (MB/sec)	Best Case Deep Arch I/O (MB/sec)
ASTER	EDC	1,055	583.8	3.4	5.7	2.4	0.8
MODIS	EDC	4,920	1,051.0	28.8	63.7	35.0	12.6
DAO	GSFC	2	13,680.0	0.9	1.8	0.9	0.3
LIS	GSFC	2	2.0	0.1	0.2	0.1	0.0
MODIS	GSFC	19,328	4,712.9	125.5	243.8	118.3	13.8
DFA/MR	JPL	114	42.2	0.1	0.1	0.1	0.0
SWS	JPL	61	45.7	0.1	0.3	0.1	0.0
CERES(AM)	LaRC	103	1,826.6	3.8	7.6	3.8	3.3
CERES(TRMM)	LaRC	103	894.7	1.5	3.0	1.5	1.0
MISR	LaRC	566	3,299.0	18.6	26.5	8.0	2.7
MOPITT	LaRC	4	9.4	0.1	0.2	0.1	0.0
SAGE	LaRC	1	3.4	0.0	0.0	0.0	0.0
MODIS	NSIDC	1,705	14.8	0.4	0.9	0.4	0.2

Dynamic Modeling Analysis

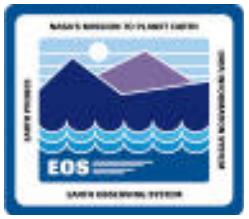

- Event driven simulation implemented using BONeS
  - Models execution of each PGE
  - Models archiving of each granule
  - Models each user pull
- Driven by technical baseline and system design
- Outputs include
  - Resource utilization versus time
  - Queue depth over time for each resource
  - Time-averaged resource utilizations
  - Elapsed time for events (e.g., PGE turn-around times)

Dynamic Modeling Status

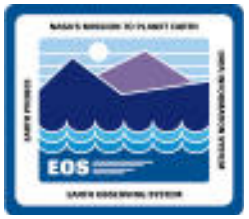

- Now using February 1996 Technical Baseline to establish push load
- Initial runs were done without MODIS Level 3s
  - Baseline simulations for EDC (ASTER), LaRC, NSIDC, JPL
  - Failover simulations for EDC (ASTER), LaRC
  - 2X, 4X, and 10X user pull simulations
- Currently modeling execution of tile-oriented MODIS Level 3s
  - Tile-oriented PGEs execute in batches
  - Results are being reviewed from full system baseline simulation including MODIS Level 3s

Dynamic Modeling Results — CERES Turnaround Times

ID	Executions	CPU (Minutes)	Queuing (Minutes)			Staging (Minutes)			Turnaround (Minutes)		
	Total	Avg	Min	Avg	Max	Min	Avg	Max	Min	Avg	Max
1aT	21.0	3.9	0.0	0.0	0.0	0.0	0.0	0.0	3.9	3.9	3.9
1aA	21.0	2.5	0.0	0.0	0.0	0.0	0.0	0.0	2.5	2.5	2.5
1bA	21.0	2.5	0.0	0.0	0.0	0.0	0.0	0.0	2.5	2.5	2.5
2aT	21.0	0.9	0.0	0.0	0.0	0.0	0.0	0.0	0.9	0.9	0.9
2aA	21.0	0.9	0.0	0.0	0.0	0.0	0.0	0.0	0.9	0.9	0.9
2bT	1.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.1	0.1
2bA	1.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.1	0.1
3aT	1.0	8.8	0.0	0.0	0.0	0.0	0.0	0.0	8.8	8.8	8.8
3aA	1.0	5.7	0.0	0.0	0.0	0.0	0.0	0.0	5.7	5.7	5.7
3bTA	1.0	17.5	0.0	0.0	0.0	0.0	0.0	0.0	17.5	17.5	17.5
4aF	504.0	13.3	0.0	1.9	12.5	0.8	18.7	60.9	14.2	33.9	86.3
4bAF	504.0	417.0	0.0	0.0	0.0	5.6	43.2	63.8	422.6	460.1	480.8
5aF	504.0	494.9	0.0	79.2	490.0	0.0	0.0	0.0	494.9	574.1	984.9
5cAF	504.0	323.9	0.0	0.0	0.0	0.0	0.0	0.0	323.9	323.9	323.9
6aT	504.0	0.9	0.0	0.5	2.3	0.0	0.0	0.0	0.9	1.4	3.2
6aA	504.0	0.6	0.0	0.0	0.0	0.0	0.0	0.0	0.6	0.6	0.6
6cT	1.0	13.8	0.0	0.0	0.0	160.6	160.6	160.6	174.4	174.4	174.4
6cA	1.0	13.8	0.0	0.0	0.0	64.9	64.9	64.9	78.7	78.7	78.7
7aT	1.0	756.0	0.0	0.0	0.0	183.3	183.3	183.3	939.3	939.3	939.3
7aA	1.0	494.8	0.0	0.0	0.0	144.7	144.7	144.7	639.5	639.5	639.5
7c	1.0	989.7	0.0	0.0	0.0	84.0	84.0	84.0	1073.7	1073.7	1073.7
8aT	1.0	42.0	0.0	0.0	0.0	0.0	0.0	0.0	42.0	42.0	42.0
8aA	1.0	27.5	0.0	0.0	0.0	0.0	0.0	0.0	27.5	27.5	27.5

Dynamic Modeling Results — LaRC CPU Usage

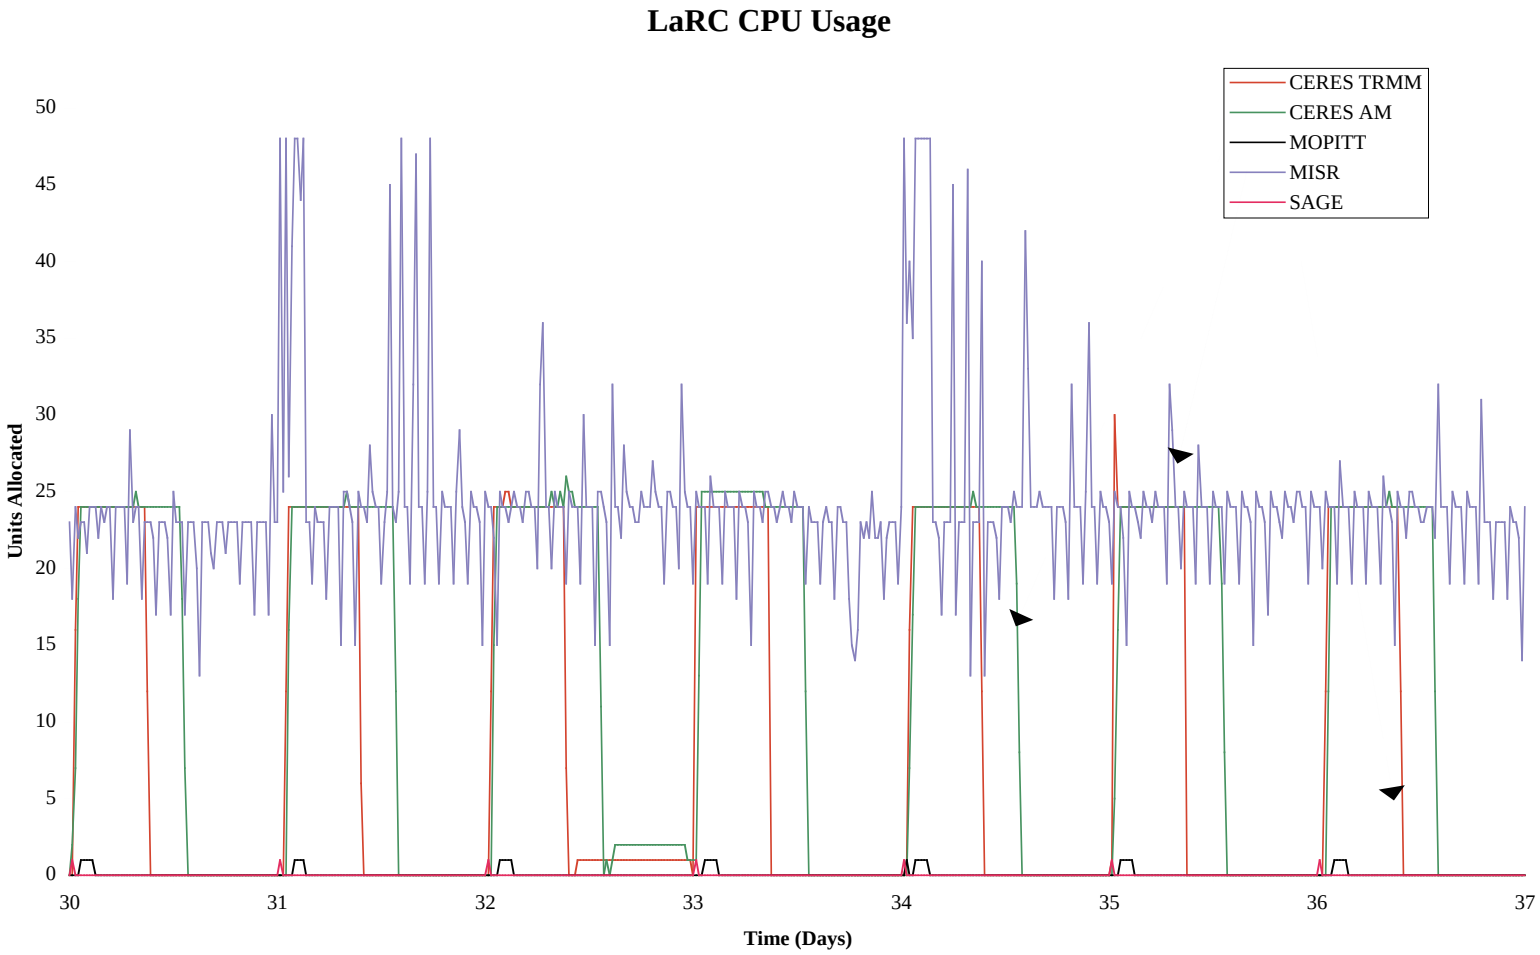

Dynamic Model Results — LaRC Execution Profile

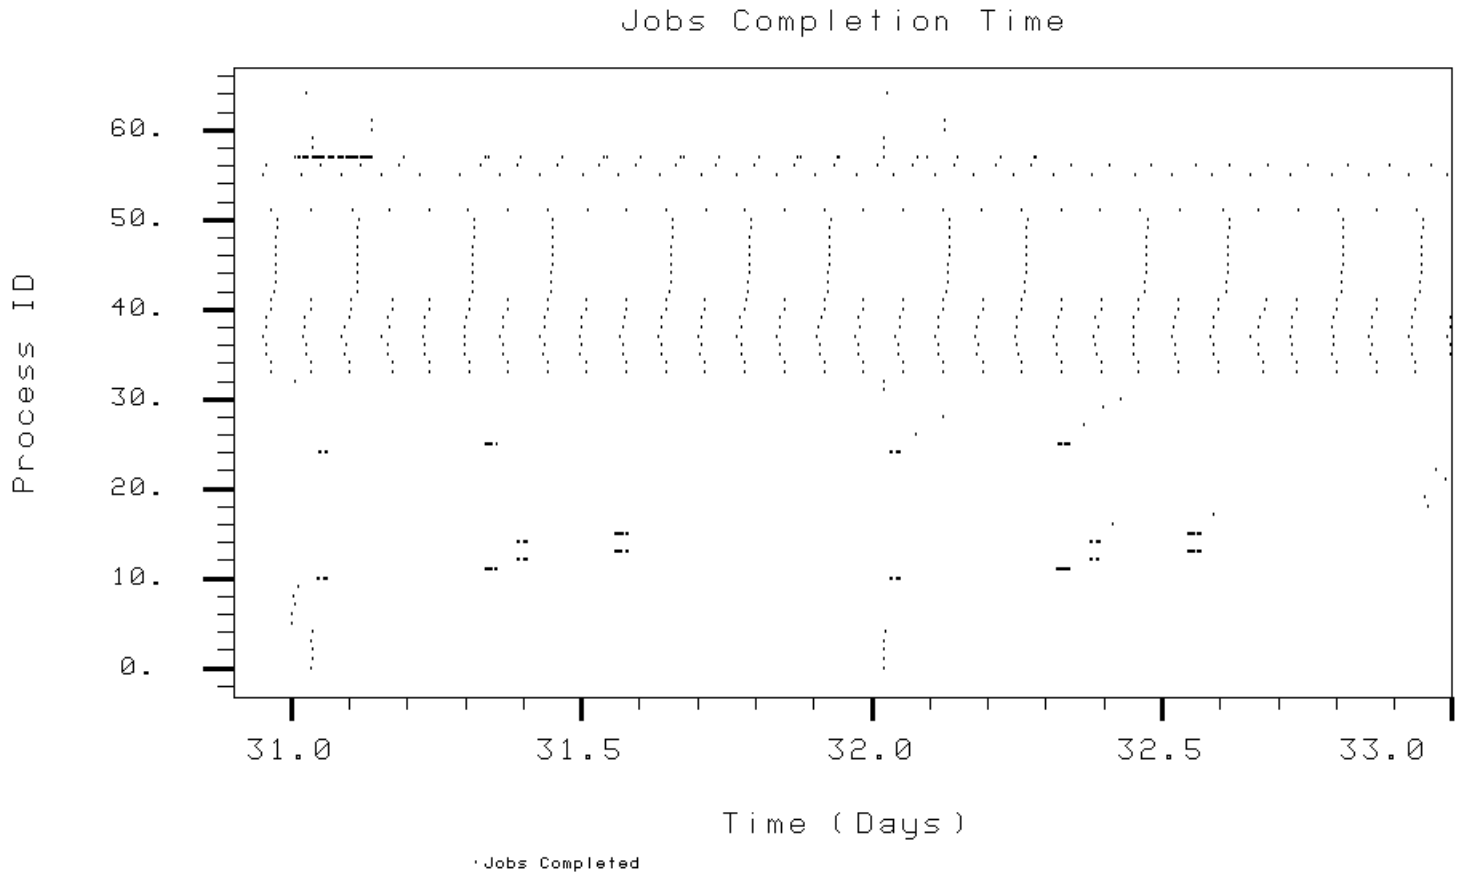

Memory and I/O Survey

- The objective of the survey was to gather the best possible data stating memory requirements and disk I/O characteristics for each PGE, to support system sizing

- The responses received to date have been minimal:
  - ASTER provided data for essentially all of their PGEs
  - CERES, MODIS, MISR, and DAO data unavailable at this time

- For the CDR design, assumptions were made
  - 128 MB of RAM per processor
  - 8 MB/sec per SCSI-2 channel/controller

- Assumptions will be validated via SSI&T and benchmarking

Other Considerations and Adjustments

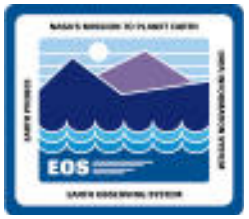

- Additions of CPU for overheads (e.g., Network I/O)
- Addition of resources to meet RMA requirements
- Re-Use and/or upgrade of existing equipment
- Configuration of disk groupings for striping
- Extrapolation from the old baseline to the new baseline where modeling results are not yet available
- Rounding up to the next configurable increment

SPRHW Specification

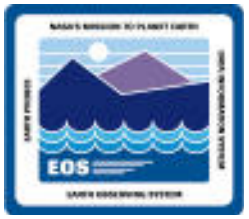

- Top-Level Summary
- Detailed System Specifications
- Specifications by Component Type
  - CPU
  - Memory
  - Disk
  - Network
  - Enclosure

SPRHW Top-Level Summary

Site	String	System	Epoch K (3Q99) Derated Processing [MF]	RAM [MB]	RAM [interleave]	Disk Channels [N]	Staging I/O [MB/s]	Processing I/O [MB/s]	Net Disk Space [GB]
EDC	**AI&T**	-1	**1,375**	**1,024**	**8**	**10**	**0.0**	**0.0**	**720**
	ASTER	-4	**1,375**	**1,024**	**4**	**2**	**3.9**	**5.3**	**110**
	MODIS	-5	1,375	1,024	8	10	32.4	24.1	720
		-6	1,375	1,024	8	10	32.4	24.1	720
		Total	**2,750**	**2,048**	**16**	**20**	**64.8**	**48.2**	**1,440**
EDC	**All**	**All**	**5,500**	**4,096**	**28**	**32**	**68.7**	**53.5**	**2,270**
GSFC	AI&T	-1							
		-9	1,100	1,024	8	8	0.0	0.0	247
		-10							
		Total	**1,100**	**1,024**	**8**	**8**	**0.0**	**0.0**	**247**
	LIS & COLOR	-4	**75**	**512**	**2**	**2**	**0.3**	**0.3**	**68**
	MODIS	-1	2,475	2,048	8	10	8.0	60.4	247
		-5	2,475	2,048	8	10	8.0	60.4	247
		-6	2,475	2,048	8	10	8.0	60.4	247
		-8	2,475	2,048	8	10	8.0	60.4	247
		-11							
		-12							
		Total	**9,900**	**8,192**	**32**	**40**	**32**	**242**	**989**
GSFC	**All**	**All**	**11,075**	**9,728**	**42**	**50**	**32.3**	**241.9**	**1,304**
JPL	**AI&T**	-1	**137**	**128**	**1**	**1**	**0.0**	**0.0**	**17**
	DFA/MR & SWS	-2	**825**	**512**	**2**	**1**	**1.8**	**1.8**	**17**
JPL	**All**	**All**	**962**	**640**	**3**	**2**	**1.8**	**1.8**	**34**
LaRC	AI&T	-1	1,925	2,048	8	4	0.0	0.0	288
		-13							
		Total	**1,925**	**2,048**	**8**	**4**	**0.0**	**0.0**	**288**
	CERES TRMM	-5	1,080	2,048	8	2	1.8	2.3	69
		-6	1,620	2,048	8	2	1.8	2.3	103
		Total	**2,700**	**4,096**	**16**	**4**	**3.6**	**4.6**	**172**
	CERES AM-1	-8	2,200	2,048	8	3	3.5	3.8	432
		-11	2,200	2,048	8	3	3.5	3.8	432
		Total	**4,400**	**4,096**	**16**	**6**	**7.0**	**7.6**	**864**
	MISR	-9	3,300	2,048	8	4	5.4	18.6	288
		-10	3,300	2,048	8	4	5.4	18.6	288
		-12							
		Total	**6,600**	**4,096**	**16**	**8**	**10.8**	**37.2**	**576**
LARC	**All**	**All**	**15,625**	**14,336**	**56**	**22**	**21.4**	**49.4**	**1,900**
NSIDC	**AI&T**	-1	**137**	**128**	**1**	**1**	**0.0**	**0.0**	**17**
	MODIS	-2	**550**	**512**	**2**	**1**	**4.0**	**4.0**	**17**
NSIDC	**All**	**All**	**687**	**640**	**3**	**2**	**4.0**	**4.0**	**34**
All	**All**	**All**	**33,849**	**29,440**	**132**	**108**	**128.2**	**350.6**	**5,542**

SPRHW Detail (Example)

C 3Q97	G 3Q98	K 3Q99
SPRHW-EDC-4	**SPRHW-EDC-4**	**SPRHW-EDC-4**
Function: EDC AI&T	Function: ASTER	Function: ASTER
Cabinet: Power Challenge XL	Cabinet: Power Challenge XL	Cabinet: Power Challenge XL
Console: Character	Console: Character	Console: Character
CPU: 6 x 275 MHz R10000	CPU: 6 x 275 MHz R10000	CPU: 10 x 275 MHz R10000
RAM: 1 GB/4-way interleaved	RAM: 1 GB/4-way interleaved	RAM: 1 GB/4-way interleaved
IO4: Two	IO4: Two	IO4: Two
HIO-1 (1,1): FDDI	HIO-1 (1,1): FDDI	HIO-1 (1,1): FDDI
HIO-2 (1,2): SCSI	HIO-2 (1,2): SCSI	HIO-2 (1,2): SCSI
HIO-3 (2,1): HiPPI	HIO-3 (2,1): HiPPI	HIO-3 (2,1): HiPPI
HIO-4 (2,2): Unused	HIO-4 (2,2): Unused	HIO-4 (2,2): Unused
SCSI-0 (1,0,1): CD-ROM	SCSI-0 (1,0,1): CD-ROM	SCSI-0 (1,0,1): CD-ROM
SCSI-1 (1,0,2): Two 4.3 GB Internal Disks	SCSI-1 (1,0,2): Two 4.3 GB Internal Disks	SCSI-1 (1,0,2): Two 4.3 GB Internal Disks
SCSI-2 (1,2,1): RAID-1 SP1	SCSI-2 (1,2,1): RAID-1 SP1	SCSI-2 (1,2,1): RAID-1 SP1
SCSI-3 (1,2,2): RAID-1 SP2	SCSI-3 (1,2,2): RAID-1 SP2	SCSI-3 (1,2,2): RAID-1 SP2
SCSI-4 (1,2,3): Unused	SCSI-4 (1,2,3): Unused	SCSI-4 (1,2,3): Unused
RAID-1: 10 x 9 GB RAID 5 (Cabinet 1)	RAID-1: 10 x 9 GB RAID 5 (Cabinet 1)	RAID-1: 15 x 9 GB RAID 5 (Cabinet 1)

SPRHW Processors

- SGI Power Challenge Processors
  - R10000
    - For new machines purchased for Release B
    - Currently shipping 200 MHz chips; 275 MHz chips announced
    - Two floating point operations per clock cycle
    - Two or four processors per board; up to 36 processors per system
  - R8000
    - Retained from Release A for CERES TRMM processing; other Release A R8000s traded in for R10000 processors
    - Retaining only 90 MHz processors
    - Four floating point operations per clock cycle
    - Two processors per board; up to 18 processors per system

SPRHW Processors — (Continued)

- SGI Power Challenge Processors
  - R4600
    - Retained from Release A in one system to support LIS and COLOR
    - Retaining only 150 MHz processors
    - One floating point operation per two clock cycles
    - One, two, or four processors per board; up to 36 processors per system
- Number of CPUs per system configured to satisfy DAAC and instrument requirements
  - Up to 20 CPUs per system at Epochs C and G (Initial purchase)
  - Up to 24 CPUs per system at Epoch K (Second purchase)
  - Some small systems configured at NSIDC, JPL, GSFC (LIS & COLOR)

Random Access Memory

- Approximately 128 MB per processor:
  - 128 MB for uniprocessors (NSIDC and JPL AI&T/AQA systems)
  - 512 MB for 4 processors (NSIDC, JPL, and LIS/COLOR systems)
  - One GB for 5 to 12 processors (EDC, GSFC AI&T systems)
  - Two GB for 13 or more processors (LaRC, GSFC systems)

- Memory Interleaving
  - Smaller systems 1 or 2 way interleaved
  - 1 GB x 4 or 8 way (Depending on I/O expected on system)
  - 2 GB x 8 way

I/O Subsystems

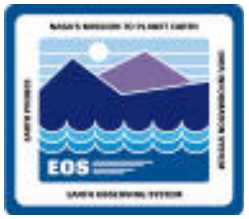

- The Challenge architecture supports up to 6 I/O subsystems (IO4 cards) per system
- Each IO4 supports up to 320 MB/s
- Each IO4 provides two FWD SCSI-2 channels and two HIO ports
- Each HIO can support
  - A card with three FWD SCSI-2 channels, or
  - A HiPPI connection, or
  - A FDDI connection
- The number of IO4 cards per system is driven by the number of connections (HiPPI + FDDI + SCSI-2) required

Internal Disk Storage

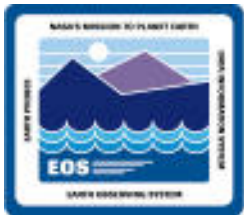

- Used for
  - swap space
  - operating system
  - COTS and ECS software
- Sized as four times the RAM size, plus two GB
- Configured as one, two, or three 4.3 or 9 GB disks, on a single SCSI-2 channel

External Disk Arrays

- Each system's external array is sized according to the I/O rates and storage requirements for its intended instrument:
  - Some instruments have a high "size to rate" ratio (CERES AM-1); they need many disks but few controllers/channels
  - Some instruments have a low "size to rate" ratio (MODIS at GSFC); they require more controllers and fewer disks/channels

- Arrays are specified as SGI RAID 5 (SCSI-2 based)
  - Waiting to see new Fibre Channel offerings under SGI "Gold Seal" program

Why Use RAID?

- Recoverability from a single disk failure
  - RAID 1 can sustain the failure of one disk per pair without data loss
  - RAID 3 and 5 can sustain the failure of one disk per group without data loss
  - With a single disk failure, the array stays on line with degraded performance
  - The SGI RAID units can replace/rebuild a single failed disk automatically using a hot spare
  - At Epoch K, we will have 260 - 315 drives in SPRHW at each of EDC, GSFC, and LaRC
    - Predicted MTBF per drive is ~300,000 hours
    - MTBF of any drive in SPRHW becomes approximately seven weeks
    - With striped filesystems required for Release B performance, loss of a single drive may require rebuilding a filesystem consisting of 16 (or more) drives (68 to 144 GB)

Why Use RAID? (Continued)

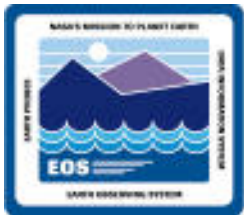

- Without RAID, limited to 7 disks per SCSI-2 channel
  - Most systems currently configured with 8 to 16 data disks per channel; would need ~20% more SCSI-2 channels, totaled across all SPRHW systems
  - Number of SCSI-2 HIO and IO4 boards would increase

- The ECS cost ratio for RAID 3 or 5 to Non-RAID Is 1.37:1.00
  - Based on a quotation for 500 GB of net storage for both approaches received from SGI on 2/28/96
  - Does not take into account additional costs of non-RAID solution for additional IO4 and SCSI-2 HIO cards

Network Interfaces

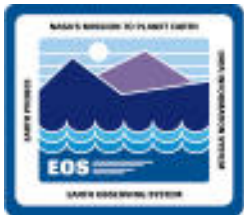

- **HiPPI**
  - **At LaRC, GSFC, and EDC, each SPRHW system (except LIS/ COLOR) will have one HiPPI interface**
- **FDDI**
  - **Each SPRHW system will have one FDDI interface**

Failure Recovery

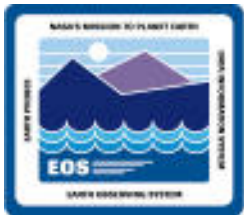

What Fails, and for How Long?

- **Science Software - Failure could be extended, but is likely to have limited impact**
- **SPRHW Hardware**
  - **Redundant components (CPU, Memory) - On call vendor maintenance should limit outage to a single shift; system may function in degraded mode**
  - **Single points of failure - On call vendor maintenance should limit outage to a single shift, but system will be unavailable during duration of failure**
- **Other subsystems/external systems - Failure could be extended, with broad impact on first-time processing, but possibly less impact on reprocessing**

Failure Recovery

Single System Failure Recovery Scenarios at Epoch K (3Q99)

- MISR, CERES, MODIS/EDC:
  - With sizing for 2X processing spread across two machines for each instrument, sufficient capacity remains to continue first-time processing
  - Reprocessing builds a backlog until repair or until AI&T system is re-configured for production
  - Reprocessing is recovered only if the AI&T resources are used temporarily to augment the production resources
- MODIS/GSFC
  - With sizing for 2X processing spread across four machines, sufficient capacity remains to continue first-time processing
  - Reprocessing continues in degraded mode, building a backlog, until repair or until AI&T system is re-configured for production
  - Reprocessing is recovered only if the AI&T resources are used temporarily to augment the production resources

Failure Recovery

Single System Failure Recovery Scenarios at Epoch K (continued)

- **EDC (ASTER), JPL, NSIDC:**
  - **First-time processing and reprocessing builds a backlog until repair or until AI&T system is re-configured for production**
  - **Before repairs,**
    - **AI&T system allows EDC(ASTER) to recover first-time processing, then keep up**
    - **AI&T systems only allow JPL and NSIDC to keep up with first-time processing (no backlog work-off)**
  - **After repairs, reprocessing is recovered only if the AI&T resources are used temporarily to augment the production resources**

Expandability

- Is there margin above requirements?
- How far can the planned boxes be expanded?
  - Up to 36 CPUs per system
  - Up to 16 GB of RAM per system
  - Up to 40 SCSI-2 channels per system
- Where will upgrades take us?
  - Faster CPUs
  - New disk technologies (Fibre Channel)
- Expansion by adding boxes

Expandability

DAAC	Percent Expandability By Resource		
	CPU	RAM	Disk I/O
EDC	260	1500	275
GSFC	125	789	317
JPL	100	700	1000
LaRC	79	700	1172
NSIDC	200	700	1000

Design Validation

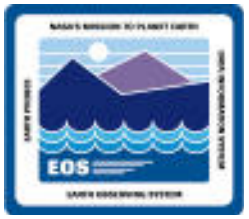

- SPRHW Benchmarking
  - Effort planned for April through July to support Release B procurement
  - Objectives are to validate design assumptions
    - Big Data Service (BDS) throughput and load
    - Filesystem throughput
    - MODIS performance calibration

- Ir1 SSI&T
  - Efforts planned through June
  - Performance data will validate AHWGP inputs and refine memory requirements

AITHW

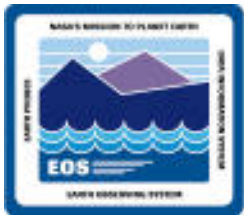

Function:

The function of AITHW is to support the integration and test of science software at the DAAC. AITHW provides tools (code management, debugging, performance) for software integration and test, and seats (development stations) for the I&T team. Remote access to the AI&T tools is also provided to the instrument teams.

Specification:

At each processing DAAC, AITHW provides a tools server (a Sun 20/50 with 128 MB of RAM and 4 GB of disk) and a number of developer's stations (Sun 20/50 workstations and/or NCD X-terminals). A target environment (SGI compute platform) for AI&T is provided in the sizing of SPRHW.

AQAHW

Function:

The function of AQAHW is to provide the DAAC with resources to perform non-science Quality Assurance testing.

Specification:

At each processing DAAC, AQAHW is provided as an SGI visualization workstation (SGI Indigo 2 IMPACT 10000 workstation, with 128 MB RAM and approximately 17 GB of disk space).